



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re U.S. Patent Application of)

FUJIBAYASHI)

Art Unit 2184

Application Number: 10/802,003)

Filed: March 17, 2004)

For: HEARTBEAT APPARATUS VIA REMOTE MIRRORING))
MIRRORING LINK ON MULTI-SITE AND)
METHOD OF USING THE SAME)

Attorney Docket No. HITC.0004)

Honorable Assistant Commissioner
for Patents
Washington, D.C. 20231

PETITION TO MAKE SPECIAL UNDER 37 C.F.R. § 1.102(d)

FOR ACCELERATED EXAMINATION

Sir:

Pursuant to 37 C.F.R. § 1.102(d), Applicant respectfully requests that the application be examined on the merits in conjunction with the pre-examination search results, the detailed discussion of the relevance of the results and amendments as filed concurrently.


Substantive consideration of the claims is respectfully solicited. Should there be any outstanding issues requiring discussion that would further the prosecution and allowance of the above-captioned application, the Examiner is invited to contact the Applicant's undersigned representative at the address and telephone number indicated below.

Respectfully submitted,

REED SMITH LLP

3110 Fairview Park Drive, Suite 1400
Falls Church, Virginia 22042
(703) 641-4200
July 1, 2005
SPF/JCM/JT

Stanley P. Fisher
Registration Number 24,344


Juan Carlos A. Marquez
Registration Number 34,072



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re U.S. Patent Application of)
FUJIBAYASHI) Art Unit 2184
Application Number: 10/802,003)
Filed: March 17, 2004)
For: HEARTBEAT APPARATUS VIA REMOTE MIRRORING)
MIRRORING LINK ON MULTI-SITE AND)
METHOD OF USING THE SAME)
Attorney Docket No. HITC.0004)
Honorable Assistant Commissioner
for Patents
Washington, D.C. 20231

STATEMENTS & PRE-EXAMINATION SEARCH REPORT
SUPPLEMENTAL TO
THE PETITION TO MAKE SPECIAL

Sir:

Pursuant to 37 C.F.R. §§ 1.102 and MPEP 708.02 VIII, Applicant hereby submits that (1) all claims of record are directed to a single invention, or if the Office determines that all the claims presented are not obviously directed to a single invention, will make an election without traverse as a prerequisite to the grant of special status; (2) a pre-examination search has been conducted according to the following field of search; (3) copies of each reference deemed most closely related to the subject matter encompassed by the claims are enclosed; and (4) a detailed discussion of the references pointing out how the claimed subject matter is patentable over the references is also enclosed herewith.

FIELD OF THE SEARCH

The field of search was directed to Class 711, subclasses 112, 114, 203, and 209 (U.S. & Foreign). Additionally, a computer database search was conducted on the U.S.P.T.O. systems EAST and WEST for U.S. and foreign patents; a keyword search was conducted in Class 711, subclasses 112, 113, 114, 147, 148, 161 and 162; Class 714, subclasses 5, 6, 7, 703, 718, 770 and 805; and a literature search was also conducted using commercial

databases for non-patent documents. Examiner Jen Tsai in Group Art Unit 2186 was consulted in confirming the field of search.

The search was directed towards a method and apparatus for encrypted remote copy for secure data backup and restoration. In particular, the search was directed towards claims 1-16 of U.S. Application Number 10/802,003. All pending claims depend from claims 1 and 11. Claim 1 is representative and sets forth a method of checking status of a first site, comprising monitoring I/O activity from a first host to a first storage system at the first site, determining status of the first host based on the I/O activity, and sending the status of the first host from the first storage system to the second storage system via the remote copy link. Claim 11 is similar and sets forth a related system implementing the above method.

With reference to the disclosure, FIG. 1A illustrates an apparatus 100 including a primary host group 101, a secondary host group 102, a primary SAN ("storage area network") 120, a secondary SAN 121, and a plurality of interchange links. See section [0041]. FIG. 1F illustrates a logical sublayer of the apparatus 100 wherein a storage system 110 includes an alert function block 116 and a storage system 111 includes an alert function block 117. See section [0054]. Each alert function block (116, 117) includes an alert engine 128 and a monitor engine 129. See section [0053]. During operation, the monitor engine 129 of the alert function block 116 sends notification signals to the alert engine 128 of the alert function block 117. See section [0054]. Likewise, the monitor engine 129 of the alert function block 117 sends notification signals to the alert engine 128 of the alert function block 116. Id. FIG. 2A illustrates a heartbeat apparatus via a remote monitoring link with a three-site multi-hoop configuration. See section [0055]. A primary host group 201 is connected with the SAN 120 through data interchange links 225.1 to 225.N. Id. A third host group 203 is not mandatory for generation operation, but maintains operation and prevents catastrophic loss. See section [0059].

With reference to the claims, a method and system check status of a first site by monitoring I/O activity from a first host to a first storage system at the first site, determining status of the first host based on the I/O activity, and sending the status of the first host from the first storage system to the second storage system via the remote copy link. (See Conclusion paragraph for detailed references to drawings and specification).

LIST OF RELEVANT REFERENCES

The search revealed the following U.S. patents, which are listed for convenience:

<u>U.S. Patent Number</u>	<u>Inventor(s)</u>
6,643,795	Sicola et al.
6,671,705	Duprey et al.
6,732,243	Busser et al.
6,834,326	Wang et
<u>U.S. Publication Number</u>	<u>Inventor(s)</u>
2004/0153841	Beck
2005/0080895	Cook et al.

Discussion of References:

U.S. Patent No. 6,643,795 to **Sicola** et al. discloses a controller-based bi-directional remote copy system with a storage site failover capability. FIG. 2 illustrates a switched dual fabric, disaster-tolerant storage system 100. See col. 7, ln. 16-30. Data storage sites 218 and 219 each respectively comprise two hosts 101/101A and 102/102A, and two storage array controllers 201/202 and 211/212 connected to storage arrays 203 and 213, respectively. Id. FIG. 4 illustrates a 'remote copy set' operation, wherein the system 100 views volumes (or LUNs) on a local array as being paired with counterpart volumes on a remote array. See col. 9, ln. 1-12. When a local host computer 101, for example, requests a storage array I/O operation, the local array controller (initiator) 301, presents a local volume that is part of the remote copy set to the local host. Id. The host 101 performs writes to the local volume on the local array 203, which copies the incoming write data to the remote volume on the target array 213. Id. One port of each controller is the "host" port that will serve LUN(s) to the local host 101/102; the other port of each controller is the "remote copy" port used for disaster tolerant backup. See col. 8, ln. 53-67. The mirroring of data for backup purposes is the basis for RAID ("Redundant Array of Independent [or Inexpensive] Disks") and the host computer sees multiple volumes and the data replication function is performed by the controller. Id. However, **Sicola** only mirrors data. **Sicola's** storage array controller 201/202 connected to storage array 203 does not monitor I/O activity from a first host to a first storage system at the first site to determine a related status of the first host. As such, **Sicola** does not

“monitor I/O activity from the first host 101 to the first storage system 110 and determine status of the first host 101 based on the I/O activity on the first host 101” as recited in claim 1, or provide “a first storage system 110 configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system 111 via the remote copy link” as recited in claim 11

U.S. Patent No. 6,671,705 to **Duprey** et al. discloses a remote mirroring system that stores information in a log and uses the information from the log to quickly resynchronize slave images following failure. FIG. 3 illustrates components of a SP (“storage port”), i.e. a primary SP 204 and a secondary SP 208, for operation in a master storage unit 130. See col. 6, ln. 62 to col. 7, ln. 29. An automatic backup/restore logic 306 automatically restores a write cache 304 from a Disk Array 206 upon detecting a failure of the master storage unit 130 and restores the write cache 304 from the Disk Array 206 when the SP 204 recovers from the failure. *Id.* Redundant copies of the information are preferably maintained within the Disk Array 206 in case of a partial disk array failure. *Id.* In order to retrieve mirror attributes (specifically, a copy of the mirror attributes for each image in the mirror), the administrator sends a request to any storage unit in the mirror, and the receiving storage unit returns a copy of the mirror attributes for each image in the mirror (or a set of error codes for any irretrievable image) to the administrator. See col. 15, ln. 25-34. The ability to treat a group of LUs (“logical units”) as a single entity simplifies the host administrator's task of managing a remote mirror for a host volume aggregated from one or more LUs. See col. 7, ln. 43-53. However, **Duprey**'s mirror attributes and storage port 204 do not monitor I/O activity from a first host to a first storage system at the first site to determine a related status of the first host. As such, **Duprey** does not “monitor I/O activity from the first host 101 to the first storage system 110 and determine status of the first host 101 based on the I/O activity on the first host 101” as recited in claim 1, or provide “a first storage system 110 configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system 111 via the remote copy link” as recited in claim 11

U.S. Patent No. 6,732,243 to **Busser** et al. discusses data mirroring in a redundant array of inexpensive disks (“RAID”) using shared buses. FIG. 2 illustrates a network storage apparatus 100 including one or more controller management modules (CMMs), e.g. CMM-A 104 and CMM-B 108. See col. 6, ln. 23-546. A passive bus backplane 116 has a first data bus 120, a second data bus 124, a third data bus 128, and a fourth data bus 132. Id. FIG. 5 illustrates a failover reset link 240 present between the CMM-A 104 and the CMM-B 108. See col. 11, ln. 30-53. The failover reset link 240 is used for communicating a failure of one of the CMMs 104, 108. Id. Thus, if a problem is detected in a heartbeat signal, the CMM 104, 108 can send a signal over the failover reset link 240 to terminate the operation of the other CMM. Id. However, **Busser**’s failover reset mechanism, storage controller management modules 104, 108 do not monitor I/O activity from a first host to a first storage system at the first site to determine a related status of the first host. As such, **Busser** does not “monitor I/O activity from the first host 101 to the first storage system 110 and determine status of the first host 101 based on the I/O activity on the first host 101” as recited in claim 1, or provide “a first storage system 110 configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system 111 via the remote copy link” as recited in claim 11

U.S. Patent No. 6,834,326 to **Wang** et al. discloses a RAID method with a network protocol between a controller and storage devices. FIG. 5 illustrates an architectural framework for a Switched RAID system, wherein the disk can be a Networked Attached Storage (“NAS”) disk, which is a thin server with a standard file protocol support. See col. 8, ln. 30-41. FIG. 7 illustrates an architecture for the RAID controller (host or switch) with NetRAID 720 viewed as a lower layer to NetSCSI or integrated with NetSCSI. See col. 11, ln. 11-28. The NetSCSI operates over different network transports. See col. 11, ln. 29-40. Status is provided after completion of commands, with important status codes being GOOD, CHECK CONDITION, and BUSY. See col. 16, ln. 60 to col. 17, ln. 5. An STP (“Standard Transport Protocol”) is provided as a heartbeat. See col. 23, ln. 14-25. However, **Wang**’s status implementation mechanism and storage controller does not monitor I/O activity from a first host to a first storage system at the first site to determine a related status of the first host. As such, **Wang** does not “monitor I/O activity from the first host 101 to the first storage system 110 and determine status of the first host 101 based on the I/O activity on the first

host 101” as recited in claim 1, or provide “a first storage system 110 configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system 111 via the remote copy link” as recited in claim 11

U.S. Pub. No. 2004/0153841 to **Beck** discloses a failure hierarchy in a cluster filesystem. FIG. 2 illustrates a cluster computing system formed of heterogeneous computer systems or nodes 22 . See section [0037]. The cluster includes a storage area network in which mass or secondary storage, such as disk drives 28 are connected to the nodes 22, 24, 26 via a Fibre Channel switch 30 and Fibre Channel connections 32. Id. One of the nodes, e.g., an IRIX node 22b, is a metadata server for the other nodes 22, 24, 26 in the cluster which are thus metadata clients with respect to the file system(s) for which the node 22b is a metadata server. See section [0042]. FIG. 3 illustrates a system wherein "vnode" 42 presents a file system independent set of operations on a file to the rest of the operating system. Id. FIG. 4, illustrates a vnode 52, which contains the head 53 of a chain of behaviors 54. See section [0043]. Each behavior points to a set of vnode operations 58 and a file system specific inode data structure 56. Id. Vnode operations 58a manage the distribution of the file metadata between nodes in the cluster, and in turn use vnode operations 58b to perform requested manipulations of the file metadata. Id. However, Beck’s vnode operations 58 do not monitor I/O activity from a first host to a first storage system at the first site to determine a related status of the first host. As such, **Beck** does not “monitor I/O activity from the first host 101 to the first storage system 110 and determine status of the first host 101 based on the I/O activity on the first host 101” as recited in claim 1, or provide “a first storage system 110 configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system 111 via the remote copy link” as recited in claim 11

U.S. Pub. No. 2005/0080895 to **Cook et al.** discusses remote activity monitoring of first server heartbeat information. A server clustering software uses TCP/IP based communications for transmitting communications, referred to as "heartbeat communications", between nodes in a cluster for determining if all of the nodes are operational. See section [0004]. FIG. 1 illustrates a cluster 100 including a first server 102 (which may also be called a node), and a first storage system 104, which are located at a local site 106. See section [0016]. The first server 102 is coupled to a first storage controller 108 in the first storage

system 104 with a link 110 (e.g. FCP --fibre channel protocol or parallel SCSI). Id. The cluster 100 includes a second server 112, and a second storage system 114, which are located at a remote site 116. Id. FIGS. 4A, 4B, and 4C illustrate a sequence 400 for remotely monitoring activity. See section [0026]. The sequence 400 begins with an operation 402, which comprises establishing first and second PPRC pairs using a Read From Secondary option. Id. Establishing a PPRC pair may comprise making an initial copy of the information on the primary site disk, onto the secondary site disk, and then entering a duplexed state, wherein there is replication of all update activity. Id. An operation 408 comprises sending (or attempting to copy) the first server heartbeat information from the first primary site disk 138, to the first secondary site disk 146 at the remote site 116. See section [0028]. Unless a pertinent portion of the cluster 100 is non-operational, the sending operation 408 will be successful, and the first server heartbeat information will be copied to the first secondary site disk 146. Id. However, **Cook's** first server heartbeat information is generated by the first server/host computer 102 itself (Abstract), rather than by the first storage system 104. In addition, **Cook's** first server/host computer 102 or first storage system 104 simply does not monitor I/O activity from the first host 102 to the first storage system 104. As such, **Cook** does not "monitor I/O activity from the first host 101 to the first storage system 110 and determine status of the first host 101 based on the I/O activity on the first host 101" as recited in claim 1, or provide "a first storage system 110 configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system 111 via the remote copy link" as recited in claim 11

Conclusion

Based on the results of the comprehensive prior art search as discussed above, Applicants contend that the method or system as recited in independent claims 1 and 11, especially the features of "monitoring I/O activity from the first host 101 to the first storage system 110; determining status of the first host 101 based on the I/O activity on the first host 101" as recited in claim 1, and "a first storage system 110 configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system

111 via the remote copy link” as recited in claim 11 are patentably distinct from the cited prior art references.

In a system including a first storage system 110 at a first site associated with a first host 101 and a second storage system 111 at a second site associated with a second host 102, the first storage system 110 and the second storage system 111 are coupled to each other by a remote copy link so that the second storage system 111 receives a copied data from the first storage system 110 via the remote copy link. The method for checking a status of the first site of the invention, as recited in claim 1, comprising: monitoring I/O activity from the first host 101 to the first storage system 110; determining status of the first host 101 based on the I/O activity on the first host 101; and sending the status of the first host 101 from the first storage system 110 to the second storage system 111 via the remote copy link.


The invention (for example, the embodiment depicted in Fig. 1F), as recited in claim 11, is directed to a data processing system comprising: a first storage system 110 at a first site associated with a first host 101; and a second storage system 111 at a second site associated with a second host 102. The first storage system 110 and the second storage system 111 are coupled to each other by a remote copy link so that the second storage system 111 receives a copied data from the first storage system 110 via the remote copy link. The first storage system 110 is configured to: monitor I/O requests received from the first host 101; determine a status of the first host 101 based on I/O activity from the first host 101, and send the status of the first host 101 to the second storage system 111 via the remote copy link.

In view of all the above, clear and distinct differences as discussed exist between the present invention as now claimed and the prior art references, Applicant respectfully contends that the prior art references cannot anticipate the present invention or render the present invention obvious. Rather, the present invention as a whole is distinguishable, and thereby allowable over the prior art.

Favorable consideration of this application is respectfully solicited. Should there be any outstanding issues requiring discussion that would further the prosecution and allowance of the above-captioned application, the Examiner is invited to contact the Applicant's undersigned representative at the address and telephone number indicated below.

Respectfully submitted,

Stanley P. Fisher
Registration Number 24,344



Juan Carlos A. Marquez
Registration Number 34,072

REED SMITH LLP
3110 Fairview Park Drive
Suite 1400
Falls Church, Virginia 22042
(703) 641-4200
July 1, 2005
SPF/JCM/JT